

CONTINUOUS APPLICATION AND DECOMPRESSION
OF TEST PATTERNS TO A CIRCUIT-UNDER-TEST

RELATED APPLICATION DATA

5 This application claims the benefit of U.S. Provisional Application No. 60/167,131, filed November 23, 1999, which is hereby incorporated by reference.

TECHNICAL FIELD

10 This invention relates generally to testing of integrated circuits and, more particularly, to the generation and application of test data in the form of patterns, or vectors, to scan chains within a circuit-under-test.

BACKGROUND

15 As integrated circuits are produced with greater and greater levels of circuit density, efficient testing schemes that guarantee very high fault coverage while minimizing test costs and chip area overhead have become essential. However, as the complexity of circuits continues to increase, high fault coverage of several types of fault models becomes more difficult to achieve with traditional testing paradigms. This difficulty arises for several reasons. First, larger integrated circuits have a very high and still increasing logic-to-pin ratio that creates a test data transfer bottleneck at the chip pins. Second, larger circuits require a prohibitively large volume of test data that must be then stored in external testing equipment. Third, applying the test data to a large circuit requires an increasingly long test application time. And fourth, present external testing equipment is unable to test such larger circuits at their speed of operation.

20 Integrated circuits are presently tested using a number of structured design for testability (DFT) techniques. These techniques rest on the general concept of making all or some state variables (memory elements like flip-flops and latches) directly controllable and observable. If this can be arranged, a circuit can be treated, as far as testing of combinational faults is concerned, as a combinational network. The most-often used DFT methodology is based on scan chains. It assumes that during testing, all (or almost all) memory elements are connected into one or more shift registers, as shown in the United States Patent No. 4,503,537. A circuit that has
30 been designed for test has two modes of operation: a normal mode and a test, or scan, mode. In

the normal mode, the memory elements perform their regular functions. In the scan mode, the memory elements become scan cells that are connected to form a number of shift registers called scan chains. These scan chains are used to shift a set of test patterns into the circuit and to shift out circuit, or test, responses to the test patterns. The test responses are then compared to fault-free responses to determine if the circuit-under-test (CUT) works properly.

Scan design methodology has gained widespread adoption by virtue of its simple automatic test pattern generation (ATPG) and silicon debugging capabilities. Today, ATPG software tools are so efficient that it is possible to generate test sets (a collection of test patterns) that guarantee almost complete fault coverage of several types of fault models including stuck-at, transition, path delay faults, and bridging faults. Typically, when a particular potential fault in a circuit is targeted by an ATPG tool, only a small number of scan cells, e.g., 2-5%, must be specified to detect the particular fault (deterministically specified cells). The remaining scan cells in the scan chains are filled with random binary values (randomly specified cells). This way the pattern is fully specified, more likely to detect some additional faults, and can be stored on a tester.

Because of the random fill requirement, however, the test patterns are grossly over-specified. These large test patterns require extensive tester memory to store and a considerable time to apply from the tester to a circuit-under-test. Fig. 1 is a block diagram of a conventional system 18 for testing digital circuits with scan chains. External automatic testing equipment (ATE), or tester, 20 applies a set of fully specified test patterns 22 one by one to a CUT 24 in scan mode via scan chains 26 within the circuit. The circuit is then run in normal mode using the test pattern as input, and the test response to the test pattern is stored in the scan chains. With the circuit again in scan mode, the response is then routed to the tester 20, which compares the response with a fault-free reference response 28, also one by one. For large circuits, this approach becomes infeasible because of large test set sizes and long test application times. It has been reported that the volume of test data can exceed one kilobit per single logic gate in a large design. The significant limitation of this approach is that it requires an expensive, memory-intensive tester and a long test time to test a complex circuit.

These limitations of time and storage can be overcome to some extent by adopting a built-in self-test (BIST) framework, as shown in the United States Patent No. 4,503,537. In BIST, additional on-chip circuitry is included to generate test patterns, evaluate test responses,

and control the test. In conventional logic BIST, where pseudo-random patterns are used as test patterns, 95-96% coverage of stuck-at faults can be achieved provided that test points are employed to address random-pattern resistant faults. On average, one to two test points may be required for every 1000 gates. In BIST, all responses propagating to observable outputs and the signature register have to be known. Unknown values corrupt the signature and therefore must be bounded by additional test logic. Even though pseudo-random test patterns appear to cover a significant percentage of stuck-at faults, these patterns must be supplemented by deterministic patterns that target the remaining, random pattern resistant faults. Very often the tester memory required to store the supplemental patterns in BIST exceeds 50% of the memory required in the deterministic approach described above. Another limitation of BIST is that other types of faults, such as transition or path delay faults, are not handled efficiently by pseudo-random patterns. Because of the complexity of the circuits and the limitations inherent in BIST, it is extremely difficult, if not impossible, to provide a set of specified test patterns that fully covers hard-to-test faults.

Weighted pseudo-random testing is another method that is used to address the issue of the random pattern resistant faults. In principle, this approach expands the pseudo-random test pattern generators by biasing the probabilities of the input bits so that the tests needed for hard-to-test faults are more likely to occur. In general, however, a circuit may require a very large number of sets of weights, and, for each weight set, a number of random patterns have to be applied. Thus, although the volume of test data is usually reduced in comparison to fully specified deterministic test patterns, the resultant test application time increases. Moreover, weighted pseudo-random testing still leaves a fraction of the fault list left uncovered. Details of weighted random pattern test systems and related methods can be found in a number of references including United States Patents Nos. 4,687,988; 4,801,870; 5,394,405; 5,414,716; and 5,612,963. Weighted random patterns have been primarily used as a solution to compress the test data on the tester. The generation hardware appears to be too complex to place it on the chip. Consequently, the voluminous test data is produced off-chip and must pass through relatively slow tester channels to the circuit-under-test. Effectively, the test application time can be much longer than that consumed by the conventional deterministic approach using ATPG patterns.

Several methods to compress test data before transmitting it to the circuit-under-test have been suggested. They are based on the observation that the test cubes (i.e., the arrangement of

test patterns bits as they are stored within the scan chains of a circuit-under-test) frequently feature a large number of unspecified (don't care) positions. One method, known as reseeding of linear feedback shift registers (LFSRs), was first proposed in B. Koenemann, "LFSR-Coded Test Patterns For Scan Designs," *Proc. European Test Conference*, pp. 237-242 (1991). Consider an n -bit LFSR with a fixed polynomial. Its output sequence is then completely determined by the initial seed. Thus, applying the feedback equations recursively provides a system of linear equations depending only on the seed variables. These equations can be associated with the successive positions of the LFSR output sequence. Consequently, a seed corresponding to the actual test pattern can be determined by solving the system of linear equations, where each equation represents one of the specified positions in the test pattern. Loading the resultant seed into the LFSR and subsequently clocking it will produce the desired test pattern. A disadvantage of this approach, however, is that seed, which encodes the contents of the test cube, is limited to approximately the size of the LFSR. If the test cube has more specified positions than the number of stages in LFSR, the test cube cannot be easily encoded with a seed. Another disadvantage of this approach is the time it requires. A tester cannot fill the LFSR with a seed concurrently with the LFSR generating a test pattern from the seed. Each of these acts must be done at mutually exclusive times. This makes the operation of the tester very inefficient, i.e., when the seed is serially loaded to the LFSR the scan chains do not operate; and when the loading of the scan chains takes place, the tester cannot transfer a seed to the LFSR.

Another compression method is based on reseeding of multiple polynomial LFSRs (MP-LFSRs) as proposed in S. Hellebrand et al., "Built-In Test For Circuits With Scan Based On Reseeding of Multiple Polynomial Linear Feedback Shift Registers," *IEEE Trans. On Computers*, vol. C-44, pp. 223-233 (1995). In this method, a concatenated group of test cubes is encoded with a number of bits specifying a seed and a polynomial identifier. The content of the MP-LFSR is loaded for each test group and has to be preserved during the decompression of each test cube within the group. The implementation of the decompressor involves adding extra memory elements to avoid overwriting the content of the MP-LFSR during the decompression of a group of test patterns. A similar technique has been also discussed in S. Hellebrand et al., "Pattern generation for a deterministic BIST scheme," *Proc. ICCAD*, pp. 88-94 (1995). Reseeding of MP-LFSRs was further enhanced by adopting the concept of variable-length seeds as described in J. Rajski et al., "Decompression of test data using variable-length seed LFSRs",

Proc. VLSI Test Symposium, pp. 426-433 (1995) and in J. Rajski et al., "Test Data Decompression for Multiple Scan Designs with Boundary Scan", *IEEE Trans. on Computers*, vol. C-47, pp. 1188-1200 (1998). This technique has a potential for significant improvement of test pattern encoding efficiency, even for test cubes with highly varying number of specified positions. The same documents propose decompression techniques for circuits with multiple scan chains and mechanisms to load seeds into the decompressor structure through the boundary-scan. Although this scheme significantly improves encoding capability, it still suffers from the two drawbacks noted above: seed-length limitations and mutually exclusive times for loading the seed and generating test patterns therefrom.

The above reseeding methods thus suffer from the following limitations. First, the encoding capability of reseeding is limited by the length of the LFSR. In general, it is very difficult to encode a test cube that has more specified positions than the length of the LFSR. Second, the loading of the seed and test pattern generation therefrom are done in two separate, non-overlapping phases. This results in poor utilization of the tester time.

A different attempt to reduce test application time and test data volume is described in I. Hamzaoglu et al., "Reducing Test Application Time For Full Scan Embedded Cores," *Proc. FTCS-29*, pp. 260-267 (1999). This so-called parallel-serial full scan scheme divides the scan chain into multiple partitions and shifts in the same test pattern to each scan chain through a single scan input. Clearly, a given test pattern must not contain contradictory values on corresponding cells in different chains loaded through the same input. Although partially specified test cubes may allow such operations, the performance of this scheme strongly relies on the scan chain configuration, i.e., the number of the scan chains used and the assignment of the memory elements to the scan chains. In large circuits such a mapping is unlikely to assume any desired form, and thus the solution is not easily scalable. Furthermore, a tester using this scheme must be able to handle test patterns of different scan chain lengths, a feature not common to many testers.

SUMMARY

A method according to the invention for applying test patterns to scan chains in a circuit-under-test includes providing a compressed test pattern of bits; decompressing the compressed test pattern into a decompressed test pattern of bits as the compressed test pattern is being

provided; and applying the decompressed test pattern to scan chains of the circuit-under-test. If desired, the method may further include applying the decompressed test pattern to scan chains of the circuit-under-test as the compressed test pattern is being provided.

The method may also include providing the compressed test pattern, decompressing the compressed test pattern, and applying the decompressed pattern synchronously. These acts may be performed at a same clock rate. Alternatively, the compressed test pattern may be provided at a lower clock rate and the compressed test pattern decompressed and applied at a higher clock rate. In yet another alternative, the compressed pattern may be provided and decompressed at a higher clock rate and the decompressed pattern applied at a lower clock rate.

Decompressing the compressed test pattern may comprise generating during a time period a greater number of decompressed test pattern bits than the number of compressed test pattern bits provided during the same time period. One way the greater number of bits may be generated is by providing a greater number of outputs for decompressed test pattern bits than the number of inputs to which the compressed test pattern bits are provided. Another way the greater number of bits may be generated is by generating the decompressed test pattern bits at a higher clock rate than the clock rate at which the compressed test pattern bits are provided.

Decompressing the compressed test pattern may further comprise generating each bit of the decompressed pattern by logically combining two or more bits of the compressed test pattern. This logically combining may include combining the bits with an XOR operation, an XNOR operation or a combination of the two operations.

In one embodiment of the invention, the providing and decompressing occur within the circuit-under-test. In another embodiment of the invention, the providing and decompressing occur within a tester, the tester applying the decompressed test pattern to scan chains of the circuit-under-test.

A circuit according to the invention may comprise a decompressor, circuit logic, and scan chains for testing the circuit logic. The decompressor is adapted to receive a compressed test pattern of bits and decompress the test pattern into a decompressed test pattern of bits as the compressed test pattern is being received. The scan chains are coupled to the decompressor and are adapted to receive the decompressed test pattern. The decompressor may comprise a linear finite state machine adapted to receive the compressed test pattern.

A tester according to the invention may comprise storage, a decompressor, and one or more tester channels. The storage is adapted to store a set of compressed test patterns of bits. The decompressor is coupled to the storage and adapted to receive a compressed test pattern of bits provided from the storage and to decompress the test pattern into a decompressed test pattern of bits as the compressed test pattern is being received. The tester channels are coupled to the decompressor and adapted to receive a decompressed test pattern and apply the decompressed test pattern to a circuit-under-test.

These and other aspects and features of the invention are described below with reference to the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a block diagram of a conventional system for testing digital circuits with scan chains.

Fig. 2 is a block diagram of a test system according to the invention for testing digital circuits with scan chains.

Fig. 3 is a block diagram of a second embodiment of a system according to the invention for testing digital circuits with scan chains.

Figs. 4A-B are block diagrams of a test system according to the invention include timing diagrams illustrating different possible timing relationships possible between the components of the system.

Fig. 5 is a block diagram of a decompressor according to the invention, including a linear finite state machine (LFSM) and phase shifter.

Fig. 6 shows in more detail a first embodiment of the decompressor of Fig. 5 coupled to a scan chain.

Fig. 7 shows the logical expressions for the bits stored in each scan cell in the scan chain of Fig. 5

Figs. 8A-8D illustrate alternative embodiments of the LFSM of Fig. 5.

Fig. 9 illustrates a preferred embodiment of a 32-bit LFSM

Figs. 10 illustrates an alternative embodiment of the phase shifter of Fig. 5.

Fig. 11 illustrates the use of parallel-to-serial conversion for applying a compressed test pattern to the decompressor.

Fig. 12 is a block diagram of a tester according to the invention for testing digital circuits with scan chains.

DETAILED DESCRIPTION

5 Fig. 2 is a block diagram of a system 30 according to the invention for testing digital circuits with scan chains. The system includes a tester 21 such as external automatic testing equipment (ATE) and a circuit 34 that includes as all or part of it a circuit-under-test (CUT) 24. The tester 21 provides from storage a set of compressed test patterns 32 of bits, one pattern at a time, through tester channels 40 to the circuit 34 such as an IC. A compressed pattern, as will be
10 described, contains far fewer bits than a conventional uncompressed test pattern. A compressed pattern need contain only enough information to recreate deterministically specified bits. Consequently, a compressed pattern is typically 2% to 5% of the size of a conventional test pattern and requires much less tester memory for storage than conventional patterns. As importantly, compressed test patterns require much less time to transmit from a tester to a CUT
15 24.

Unlike in the prior reseeding techniques described above, the compressed test patterns 32 are continuously provided from the tester 21 to scan chains 26 within the CUT 24 without interruption. As the compressed test pattern is being provided by the tester 21 to the input channels of a decompressor 36 within the circuit 34, the decompressor decompresses the
20 compressed pattern into a decompressed pattern of bits. The decompressed test pattern is then applied to the scan chains 26. This application is preferably done while the compressed test pattern is being provided to the circuit 34, but it need not be so. After circuit logic within the CUT 24 is clocked with a decompressed test pattern in the scan chains 26, the test response to that pattern is captured in the scan chains and transmitted to the tester 21 for comparison with the
25 compressed fault-free reference responses 41 stored therein.

In a typical configuration, the decompressor 36 has one output per scan chain 26, and there are more scan chains than input channels to the decompressor. However, as will be described, other configurations are also possible in which the decompressor outputs are fewer than or equal to the input channels. The decompressor generates in a given time period a greater
30 number of decompressed bits at its outputs than the number of compressed pattern bits it receives

during the same time period. This is the act of decompression, whereby the decompressor 36 generates a greater number of bits than are provided to it in a given time period.

To reduce the data volume of the test response and the time for sending the response to the tester, the circuit 34 can include means for compressing the test response that is read from the scan chains 26. One structure for providing such compression is one or more spatial compactors 38. The compressed test responses produced by the compactors 38 are then compared one by one with compressed reference responses 40. A fault is detected if a reference response does not match an actual response. Fig. 3 shows another structure that can be used for compressing the test response. A multiple input signature register (MISR) 42 compresses multiple test pattern responses into a signature that is then sent to the tester. There it is compared to a reference signature 44. Compacting the test response in the above ways is desirable but not necessary to the present decompression method and system.

The providing of a compressed test pattern to a circuit, its decompression into a decompressed test pattern, and the application of the decompressed test pattern to the scan chains is performed synchronously, continuously, and substantially concurrently. The rate at which each act occurs, however, can vary. All acts can be performed synchronously at a same clock rate if desired. Or the acts can be performed at different clock rates. If the acts are performed at the same clock rate, or if the compressed test patterns are provided and decompressed at a higher clock rate than at which the decompressed test patterns are applied to the scan chains, then the number of outputs of decompressor 36 and associated scan chains will exceed the number of input channels of the decompressor, as in Fig. 2. In this first case, decompression is achieved by providing more decompressor outputs than input channels. If the compressed test patterns are provided at a lower clock rate and decompressed and applied to the scan chains at a higher clock rate, then the number of outputs and associated scan chains can be the same, fewer, or greater than the number of input channels. In this second case, decompression is achieved by generating the decompressed test pattern bits at a higher clock rate than the clock rate at which the compressed test pattern bits are provided.

Fig. 4A illustrates an embodiment of the first case in which the compressed pattern is provided and decompressed at a higher clock rate and the decompressed pattern is applied synchronously to the scan chains at a lower clock rate. The tester 21 provides the bits of the compressed pattern through a tester channel 40 to an input channel 37 of the decompressor 36 at

a higher rate set by clock 0 (C0). The decompressor is clocked by clock 1 (C1) at the same rate as the tester and produces at outputs 39 the bits of the decompressed pattern at that rate. These decompressed bits, however, are applied to the scan chains 26 at a lower rate set by clock 2 (C2), which clocks the bits into the scan chains. This difference in rates is illustrated in the exemplary timing diagram in Fig. 4A (the actual difference can be much greater). Because of the difference therein, only every other output of the decompressor is written to the scan chains. But that is taken into account in the initial test pattern generation. One advantage of clocking the tester, decompressor, and scan chains as shown is that the tester requires fewer channels than the number of scan chains to provide the test pattern to the CUT 24. By clocking the tester at a higher clock rate C0, the time required to apply the compressed test pattern to the circuit 34 is significantly reduced. Another advantage is in low power applications, where the power dissipated during test mode has to be controlled. This can be done by reducing the clock rate C2 at which bits are shifted into the scan chains.

Fig. 4B illustrates an embodiment of the second case in which the compressed test pattern is provided at a lower clock rate and decompressed and applied synchronously at a higher clock rate. Here, the tester 21 provides the bits of the compressed pattern through channels 40 to the input channels 37 of the decompressor 36 at a lower rate set by clock 0 (C0). The decompressor is clocked by clock 1 (C1) at a higher rate. The decompressed bits are applied through its outputs 39 to the scan chains 26 by clock 2 (C2) at the same rate as clock 1. This difference in rates is illustrated in the exemplary timing diagram in Fig. 4B (the actual difference can be much greater). Because of the difference, the decompressor 36 reads the same bits from the tester 21 twice before they change. The decompressor, however, includes a state machine, as will be described, and its outputs change each clock cycle because its internal states change. One advantage of clocking the tester, decompressor, and scan chains as shown in Fig. 4B is that one can utilize a tester 21 that has many channels but with little memory behind them. By providing bits on more tester channels per clock cycle, the lack of memory depth is overcome and the time required for applying the compressed test pattern is reduced.

Fig. 5 is a block diagram of a decompressor according to the invention. In a preferred embodiment, decompressor 36 comprises a linear finite state machine (LFSM) 46 coupled, if desired, through its taps 48 to a phase shifter 50. The LFSM through the phase shifter provides highly linearly independent test patterns to the inputs of numerous scan chains in the CUT 24.

The LFSM can be built on the basis of the canonical forms of linear feedback shift registers, cellular automata, or transformed LFSRs that can be obtained by applying a number of m -sequence preserving transformations. The output of the LFSM is applied to the phase shifter, which ensures that the decompressed pattern bits present within each of the multiple scan chains 26 at any given time do not overlap in pattern (i.e., are out of phase).

The concept of continuous flow decompression described herein rests on the fact noted above that deterministic test patterns typically have only between 2 to 5 % of bits deterministically specified, with the remaining bits randomly filled during test pattern generation. (Test patterns with partially specified bit positions are called test cubes, an example of which appears in Table 2.) These partially specified test cubes are compressed so that the test data volume that has to be stored externally is significantly reduced. The fewer the number of specified bits in a test cube, the better is the ability to encode the information into a compressed pattern. The ability to encode test cubes into a compressed pattern is exploited by having a few decompressor input channels driving the circuit-under-test, which are viewed by the tester as virtual scan chains. The actual CUT 24, however, has its memory elements connected into a large number of real scan chains. Under these circumstances, even a low-cost tester that has few scan channels and sufficiently small memory for storing test data can drive the circuit externally.

Fig. 6 shows in more detail a first embodiment of the decompressor of Fig. 5. The LFSM is embodied in an eight stage Type 1 LFSR 52 implementing primitive polynomial $h(x) = x^8 + x^4 + x^3 + x^2 + 1$. The phase shifter 50, embodied in a number of XOR gates, drives eight scan chains 26, each eight bits long. The structure of the phase shifter is selected in such a way that a mutual separation between its output channels C0-C7 is at least eight bits, and all output channels are driven by 3-input (tap) XOR functions having the following forms:

$$\begin{aligned} C_0 &= s_4 \oplus s_3 \oplus s_1 & C_4 &= s_4 \oplus s_2 \oplus s_1 \\ C_1 &= s_7 \oplus s_6 \oplus s_5 & C_5 &= s_5 \oplus s_2 \oplus s_0 \\ C_2 &= s_7 \oplus s_3 \oplus s_2 & C_6 &= s_6 \oplus s_5 \oplus s_3 \\ C_3 &= s_6 \oplus s_1 \oplus s_0 & C_7 &= s_7 \oplus s_2 \oplus s_0 \end{aligned}$$

Table 1

where C_i is the i th output channel and s_k indicates the k th stage of the LFSR. Assume that the LFSR is fed every clock cycle through its two input channels 37a, 37b and input injectors 48a,

48b (XOR gates) to the second and the sixth stages of the register. The input variables “a” (compressed test pattern bits) received on channel 37a are labeled with even subscripts (a_0, a_2, a_4, \dots) and the variables “a” received on channel 37b are labeled with odd subscripts (a_1, a_3, a_5, \dots). Treating these external variables as Boolean, all scan cells can be conceptually filled with symbolic expressions being linear functions of input variables injected by tester 21 into the LFSR 52. Given the feedback polynomial, the phase shifter 50, the location of injectors 48a, b as well as an additional initial period of four clock cycles during which only the LFSR is supplied by test data, the contents of each scan cell within the scan chains 26 in Fig. 6 can be logically determined. Fig. 7 gives the expressions for the 64 scan cells in Fig. 6, with the scan chains numbered 0 through 7 in Fig. 6 corresponding to the scan chains C7, C1, C6, ... identified in Fig. 6. The expressions for each scan chain in Fig. 7 are listed in the order in which the information is shifted into the chain, i.e., the topmost expression represents the data shifted in first.

Assume that the decompressor 36 in Fig. 6 is to generate a test pattern based on the following partially specified test cube in Table 2 (the contents of the eight scan chains are shown here horizontally, with the leftmost column representing the information that is shifted first into the scan chains):

x x x x x x x	scan chain 0
x x x x x x x	scan chain 1
x x x x 1 1 x x	scan chain 2
x x 0 x x x 1 x	scan chain 3
x x x x 0 x x 1	scan chain 4
x x 0 x 0 x x x	scan chain 5
x x 1 x 1 x x x	scan chain 6
x x x x x x x	scan chain 7

Table 2

The variable x denotes a “don’t care” condition. Then a corresponding compressed test pattern can be determined by solving the following system of ten equations from Fig. 7 using any of a number of well-known techniques such as Gauss-Jordan elimination techniques. The selected equations correspond to the deterministically specified bits:

5

10

$$\begin{aligned}
 a_2 \oplus a_6 \oplus a_{11} &= 1 \\
 a_0 \oplus a_1 \oplus a_4 \oplus a_8 \oplus a_{13} &= 1 \\
 a_4 \oplus a_5 \oplus a_9 \oplus a_{11} &= 0 \\
 a_0 \oplus a_2 \oplus a_5 \oplus a_{12} \oplus a_{13} \oplus a_{17} \oplus a_{19} &= 1 \\
 a_1 \oplus a_2 \oplus a_4 \oplus a_5 \oplus a_6 \oplus a_8 \oplus a_{12} \oplus a_{15} &= 0 \\
 a_0 \oplus a_1 \oplus a_3 \oplus a_5 \oplus a_7 \oplus a_8 \oplus a_{10} \oplus a_{11} \oplus a_{12} \oplus a_{14} \oplus a_{18} \oplus a_{21} &= 1 \\
 a_2 \oplus a_3 \oplus a_4 \oplus a_9 \oplus a_{10} &= 0 \\
 a_0 \oplus a_1 \oplus a_2 \oplus a_6 \oplus a_7 \oplus a_8 \oplus a_{13} \oplus a_{14} &= 0 \\
 a_3 \oplus a_4 \oplus a_5 \oplus a_6 \oplus a_{10} &= 1 \\
 a_0 \oplus a_1 \oplus a_3 \oplus a_7 \oplus a_8 \oplus a_9 \oplus a_{10} \oplus a_{14} &= 1
 \end{aligned}$$

Table 3

It can be verified that the resulting seed variables a_0, a_1, a_2, a_3 and a_{13} are equal to the value of one while the remaining variables assume the value of zero. This seed will subsequently produce a fully specified test pattern in the following form (the initial specified positions are underlined):

15

20

25

1	0	1	0	0	1	0	0
1	1	0	0	0	1	0	0
1	1	1	1	<u>1</u>	<u>1</u>	1	0
0	0	0	1	<u>0</u>	<u>0</u>	<u>1</u>	1
1	0	<u>1</u>	0	0	0	<u>0</u>	<u>1</u>
1	1	<u>0</u>	1	<u>0</u>	0	0	<u>0</u>
1	1	<u>1</u>	1	<u>1</u>	1	1	1
0	1	0	0	<u>1</u>	1	0	0

Table 4

As can be observed, the achieved compression ratio (defined as the number of scan cells divided by the number of compressed pattern bits) is $64 / (2 \times 8 + 2 \times 4) \approx 2.66$. The fully specified test pattern is then compressed into a compressed pattern of bits using any of a number of known methods.

30

35

Figs. 8A-D illustrate various embodiments for the LFSM 46 of Fig. 5. Fig. 8A is a Type I LFSR 60. Fig. 8B is a Type II LFSR 62. Fig 8C is a transformed LFSR 64. And Fig. 8D is a cellular automaton 66. All of them implement primitive polynomials. Except for the cellular automaton 66, in each case the LFSM includes a number of memory elements connected in a shift register configuration. In addition, there are several feedback connections between various memory cells that uniquely determine the next state of the LFSM. The feedback connections are assimilated into the design by introducing injectors in the form of XOR gates near the destination

memory elements. The input channels 37 provide the bits of a compressed pattern to the LFSM through input injectors 48a, b. The injectors are handled similarly as the other feedback connections within the LFSM except that their sources of bits are the input channels. The input channels 37 may have multiple fan-outs driving different LFSM injectors 48 to improve the encoding efficiency.

Fig. 9 shows a preferred embodiment of a 32-bit LFSM in the form of a re-timed LFSR 68. The injectors are spaced equally so that the input variables are distributed optimally once they are injected into the LFSM. In practice, the size of the LFSM depends on the number of real scan chains in a circuit, the desired compression ratio of encoding, and on certain structural properties of the circuit-under-test.

Fig. 10 illustrates an alternative embodiment of a phase shifter 50, constructed with an array of XNOR gates rather than XOR gates. Phase shifters can be constructed with combinations of XNOR and XOR gates as well.

Fig. 11 illustrates the use of parallel-to-serial conversion for applying a compressed test pattern to the decompressor. If the input channels 37 to the decompressor 36 are fewer in number than the number of channels 40 of the tester 21, it can be advantageous to provide a parallel-to-serial converter such as registers 70 at the input to the decompressor. The registers 70 are clocked such that their contents are shifted out before the next set of bits is applied to the register from the tester 21. The continuous flow of the test patterns is thus preserved.

Fig. 12 is a block diagram of a tester 21 embodiment that includes the decompressor 36, rather than providing it in the circuit 34. The tester decompresses the test pattern internally and transmits the decompressed test pattern to the CUT 24. Such a tester has advantages where testing time is not as critical and it is preferred not to add a decompressor to each circuit-under-test. Storage requirements are still reduced because compressed test patterns (rather than full test patterns) need only be stored. In addition, in a variation of the above tester embodiment, the compactors 38 can also be included in the tester 21 rather than the circuit 34. The circuit then returns uncompressed test responses to the tester. This further simplifies the circuit's design.

The process of decompressing a test pattern will now be described in more detail, with reference to Fig. 5. The LFSM 46 starts its operation from an initial all-zero state. Assuming an n -bit LFSM and m input injectors, $\lceil n/m \rceil$ clock cycles may be used to initialize the LFSM before it starts generating bits corresponding to the actual test patterns. After initialization of the LFSM

and assuming clocks C0 and C1 are running at the same rate, a new bit is loaded in parallel into each scan chain 26 every clock cycle via the phase shifter 50. At this time, the circuit-under-test 34 is operated in the scan mode, so that the decompressed test pattern fills the scan chains 26 with 0s and 1s (and shifts out any previous test response stored there). A small number of bit positions in the scan chains, therefore, get deterministically specified values while the remaining positions are filled with random bits generated by the LFSM. The number of clock cycles for which a test pattern is shifted is determined by the length of the longest scan chain within the circuit, the number being at least as great as the number of cells in the longest scan chain. A scan-shift signal is therefore held high for all the scan chains until the longest scan chain gets the entire test pattern. The shorter scan chains in the circuit are left justified so that the first few bits that are shifted are overwritten without any loss of information.

Patterns from the LFSM may be linearly dependent. In other words, it is possible to determine various bit positions within the two-dimensional structure of multiple scan chains that are significantly correlated. This causes testability problems, as it is often not possible to provide the necessary stimulus for fault excitation to the gates driven by positions that have some form of dependency between them. Consequently, the phase shifter 50 (such as an array of XOR gates or XNOR gates) may be employed at the taps (outputs) of the LFSM to reduce linear dependencies between various bit positions within the scan chains. The XOR logic can be two-level or multi-level depending on the size of the XOR gates. Every scan chain in the CUT 24 is driven by signals that are obtained by XOR-ing a subset of taps 48 from the LFSM. These taps are determined so that the encoding efficiency of the test cubes is still preserved. In addition, the taps are selected in a manner so that all memory cells in the LFSM have approximately equal number of fan-out signals and the propagation delays are suitably optimized. Once a decompressed test pattern is completely loaded into the scan chains during test mode, the CUT 24 is switched to the normal mode of operation. The CUT then performs its normal operation under the stimulus provided by the test pattern in the scan chains. The test response of the CUT is captured in the scan chains. During the capture the LFSM is reset to the all-zero state before a new initialization cycle begins for loading the next test pattern.

Having illustrated and described the principles of the invention in exemplary embodiments, it should be apparent to those skilled in the art that the illustrative embodiments can be modified in arrangement and detail without departing from such principles. In view of the

many possible embodiments to which the principles of the invention may be applied, it should be understood that the illustrative embodiment is intended to teach these principles and is not intended to be a limitation on the scope of the invention. We therefore claim as our invention all that comes within the scope and spirit of the following claims and their equivalents.

000270" T2002960